

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 313 018 A1**

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
21.05.2003 Bulletin 2003/21

(51) Int Cl.7: **G06F 11/20, G06F 3/06**

(21) Application number: **02024809.2**

(22) Date of filing: **07.11.2002**

(84) Designated Contracting States:
**AT BE BG CH CY CZ DE DK EE ES FI FR GB GR
IE IT LI LU MC NL PT SE SK TR**
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **14.11.2001 US 332991 P
30.11.2001 US 998494**

(71) Applicant: **EMC Corporation
Hopkinton, MA 01748-9103 (US)**

(72) Inventors:
• **Halstead, Mark J.**
Waltham, Massachusetts 02451 (US)
• **Arnon, Dan**
Boston, Massachusetts 02116 (US)
• **Ofer, Adi**
Wellesley, Massachusetts 02481 (US)

(74) Representative: **Freischem, Stephan, Dipl.-Ing.**
Patentanwälte Freischem
An Gross St. Martin 2
50667 Köln (DE)

(54) **Hierarchical approach to identifying changing device characteristics**

(57) Determining device characteristics includes obtaining a first globally accessible value, if the first globally accessible value corresponds to a stored first value, obtaining device characteristics data from a relatively fast memory, if the first globally accessible value does not correspond to the stored first value, obtaining a second globally accessible value, if the second globally accessible value corresponds to a stored second value, obtaining device characteristics data from a relatively

fast memory, if the second globally accessible value does not correspond to the stored second value, obtaining device characteristics data from a relatively slow memory and updating the relatively fast memory, the stored first value, and the stored second value. The globally accessible first value may include device I/O information. The globally accessible values may be stored in global memory that is accessible to a plurality of processors.

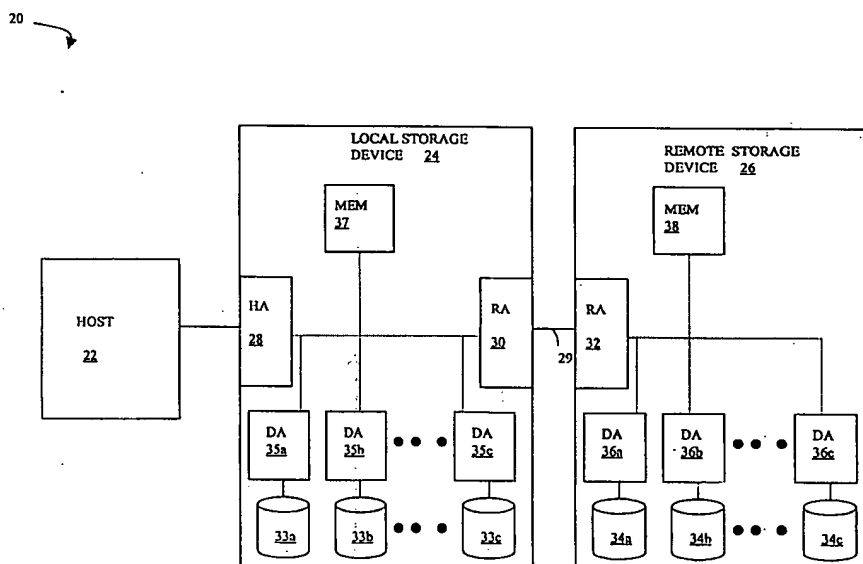


FIGURE 1

Description

Background of the Invention

1. Technical Field

[0001] This application relates to computer storage devices, and more particularly to communication between storage devices.

2. Description of Related Art

[0002] Host processor systems may store and retrieve data using a storage device containing a plurality of host interface units (host adapters), disk drives, and disk interface units (disk adapters). Such storage devices are provided, for example, by EMC Corporation of Hopkinton, Mass. and disclosed in U.S. Patent No. 5,206,939 to Yanai et al., 5,778,394 to Galtzur et al., U.S. Patent No. 5,845,147 to Vishnitzky et al., and U.S. Patent No. 5,857,208 to Ofek. The host systems access the storage device through a plurality of channels provided therewith. Host systems provide data and access control information through the channels to the storage device and the storage device provides data to the host systems also through the channels. The host systems do not address the disk drives of the storage device directly, but rather, access what appears to the host systems as a plurality of logical disk units. The logical disk units may or may not correspond to the actual disk drives. Allowing multiple host systems to access the single storage device unit allows the host systems to share data stored therein.

[0003] In some instances, it may be desirable to copy data from one storage device to another. For example, if a host writes data to a first storage device, it may be desirable to copy that data to a second storage device provided in a different location so that if a disaster occurs that renders the first storage device inoperable, the host (or another host) may resume operation using the data of the second storage device. Such a capability is provided, for example, by the Remote Data Facility (RDF) product provided by EMC Corporation of Hopkinton, Massachusetts. With RDF, a user may denote a first storage device as a master storage device and a second storage device as a slave storage device. Other incarnations of RDF may provide a peer to peer relationship between the local and remote storage devices. The host interacts directly with the local storage device, but any data changes made to the local storage device are automatically provided to a remote storage device using RDF. The local and remote storage devices may be connected by a data link, such as an ESCON link or a Fiber Channel link. The RDF functionality may be facilitated with an RDF adapter (RA) provided at each of the storage devices.

[0004] In some instances, it may be desirable to modify the RDF configuration system. However, in many

cases, such modifications require skilled technicians using special software and non-standard connections to the local storage devices. It is desirable to automate the RDF configuration modification process to allow a host to modify the RDF configuration. In addition, it is desirable that allowing dynamic modifications to RDF configuration will not effect operation of the storage device when individual devices therein need to access the dynamic configuration information.

Summary of the Invention

[0005] Object of the invention is to improve communication between storage devices. This object leads to a set of inventions which is herein described in its entirety and which is claimed in the claims of a series of parallel patent applications.

[0006] According to the present invention, dynamically creating a communication path between first and second storage devices, includes creating a connection to a source volume on the first storage device and indicating that the source volume is not ready to transmit data on the communication path, after successfully creating the connection to the source volume, creating a connection to a destination volume on the second storage device and initially indicating that portions of one of: the destination volume and the source volume do not contain valid copies of data, wherein the destination volume accepts data from the source volume, and after successfully creating the connections to the source and destination volumes, indicating that the source volume is ready to transmit data on the communication path. Dynamically creating a communication path between first and second storage devices, may also include creating at least one of: the source volume and the destination volume. Creating the connection to the source volume may include modifying a table containing configuration information for the first storage device. Creating the connection to the destination volume may include modifying a table containing configuration information for the second storage device. Dynamically creating a communication path between first and second storage devices, may further include, following unsuccessfully creating a connection to the destination volume, destroying the connection to the source volume, in which case, an error may be returned. Portions of the destination volume may be initially indicated as not containing valid data, in which case, dynamically creating a communication path between first and second storage devices, may further include, after indicating that the source volume is ready to transmit data on the communication path, initiating a background copy operation to copy data from the source volume to the destination volume. Portions of the source volume may be initially indicated as not containing valid data, in which case, dynamically creating a communication path between first and second storage devices, may further include, after indicating that the source volume is ready to transmit data on the communication

path, initiating a background copy operation to copy data from the destination volume to the source volume. The host may perform an I/O operation on a particular portion of the source volume, in which case, dynamically creating a communication path between first and second storage devices, may further include, in response to the particular portion being indicated as containing invalid data, copying data corresponding to the particular portion from the destination volume to the source volume prior to completing the I/O operation.

[0007] According further to the present invention, dynamically creating a communication path between first and second storage devices, includes creating a connection to a destination volume on the first storage device, after successfully creating the connection to the destination volume, creating a connection to a source volume on the second storage device and indicating that the source volume is not ready to transmit data on the communication path and initially indicating that portions of one of: the destination volume and the source volume do not contain valid copies of data, wherein the destination volume accepts data from the source volume and, after successfully creating the connections to the source and destination volumes, indicating that the source volume is ready to transmit data on the communication path. Dynamically creating a communication path between first and second storage devices may further include creating at least one of: the source volume and the destination volume. Creating the connection to destination volume may include modifying a table containing configuration information for the first storage device. Creating the connection to the source volume may include modifying a table containing configuration information for the second storage device. Dynamically creating a communication path between first and second storage devices may further include, following unsuccessfully creating a connection to the source volume, destroying the connection to the destination volume, in which case, an error indication may be returned. Portions of the destination volume may be initially indicated as not containing valid data. Dynamically creating a communication path between first and second storage devices may further include, after indicating that the source volume is ready to transmit data on the communication path, initiating a background copy operation to copy data from the source volume to the destination volume. Portions of the source volume may be initially indicated as not containing valid data. Dynamically creating a communication path between first and second storage devices may further include, after indicating that the source volume is ready to transmit data on the communication path, initiating a background copy operation to copy data from the destination volume to the source volume. Dynamically creating a communication path between first and second storage devices may further include, the host performing an I/O operation on a particular portion of the source volume. Dynamically creating a communication path between first and second storage

devices may further include, in response to the particular portion being indicated as containing invalid data, copying data corresponding to the particular portion from the destination volume to the source volume prior to completing the I/O operation.

[0008] According further to the present invention, a computer program product that creates a communication path between first and second storage devices, includes executable code that creates a connection to a source volume on the first storage device and indicates that the source volume is not ready to transmit data on the communication path, executable code that creates a connection to a destination volume on the second storage device and initially indicates that portions of one of: the destination volume and the source volume do not contain valid copies of data after successfully creating the connection to the source volume, wherein the destination volume accepts data from the source volume, and executable code that indicates that the source volume is ready to transmit data on the communication path after successfully creating the connections to the source and destination volumes. The computer program product may further include executable code that creates at least one of: the source volume and the destination volume. The executable code that creates the connection to the source volume modifies a table may contain configuration information for the first storage device. Executable code that creates the connection to the destination volume may modify a table containing configuration information for the second storage device. The computer program product may further include executable code that destroys the connection to the source volume following unsuccessfully creating a connection to the destination volume, in which case an error indication may be returned. The computer program product may further include executable code that causes portions of the source volume to be initially indicated as not containing valid data. The computer program product may further include executable code that initiates a background copy operation to copy data from the destination volume to the source volume after indicating that the source volume is ready to transmit data on the communication path. The computer program product may further include executable code that copies data corresponding to a requested portion from the destination volume to the source volume prior to completing an I/O operation in response to the requested portion being indicated as containing invalid data.

[0009] According further to the present invention, a computer program product that dynamically creates a communication path between first and second storage devices, includes executable code that creates a connection to a destination volume on the first storage device, executable code that creates a connection to a source volume on the second storage device and indicates that the source volume is not ready to transmit data on the communication path and initially indicates that portions of one of: the destination volume and the

source volume do not contain valid copies of data after successfully creating the connection to the destination volume, wherein the destination volume accepts data from the source volume, and executable code that indicates that the source volume is ready to transmit data on the communication path after successfully creating the connections to the source and destination volumes. The computer program product may further include executable code that creates at least one of: the source volume and the destination volume. Executable code that creates the connection to the source volume may modify a table containing configuration information for the first storage device. Executable code that creates the connection to the destination volume may modify a table containing configuration information for the second storage device. The computer program product may further include executable code that destroys the connection to the destination volume following unsuccessfully creating a connection to the source volume, in which case an error indication may be returned.

[0010] According to the present invention, reversing a communication path between a first volume on a first storage device and a second volume on a second storage device includes suspending communication between the first and second volumes while maintaining operations for other volumes of the storage devices, causing the first volume to change from a source volume to a destination volume without destroying the first volume, causing the second volume to change from a destination volume to a source volume without destroying the second volume, and resuming communication between the first and second volumes. Causing the first volume to change from a source volume to a destination volume may include modifying a table of the first storage device. Causing the second volume to change from a source volume to a destination volume may include modifying a table of the second storage device. Suspending communication may include setting the first volume to a not ready state. Resuming communication may include setting the second volume to a ready state. Reversing a communication path between a first volume on a first storage device and a second volume on a second storage device may also include returning a result indicating successfully reversing the communication path.

[0011] According further to the present invention, managing volumes on storage devices, includes receiving a command requesting reversal of a communication path between a first volume on a first storage device and a second volume on a second storage device, suspending communication between the first and second volumes while maintaining operations for other volumes of the storage devices, causing the first volume to change from a source volume to a destination volume without destroying the first volume, causing the second volume to change from a destination volume to a source volume without destroying the second volume, and resuming communication between the first and second volumes.

The command may be a single multihop, multiexecute command that causes operations to be performed on the first and second storage devices. Causing the first volume to change from a source volume to a destination volume may include modifying a table of the first storage device. Causing the second volume to change from a source volume to a destination volume may include modifying a table of the second storage device. Suspending communication may include setting the first volume to a not ready state. Resuming communication may include setting the second volume to a ready state. Managing volumes on storage devices may include returning a result indicating successfully reversing the communication path.

[0012] According further to the present invention, a computer program product that reverses a communication path between a first volume on a first storage device and a second volume on a second storage device, includes executable code that suspends communication between the first and second volumes while maintaining operations for other volumes of the storage devices, executable code that causes the first volume to change from a source volume to a destination volume without destroying the first volume, executable code that causes the second volume to change from a destination volume to a source volume without destroying the second volume, and executable code that resumes communication between the first and second volumes. Executable code that causes the first volume to change from a source volume to a destination volume may modify a table of the first storage device. Executable code that causes the second volume to change from a source volume to a destination volume may modify a table of the second storage device. Executable code that suspends communication may set the first volume to a not ready state. Executable code that resumes communication may set the second volume to a ready state. A computer program product that reverses a communication path between a first volume on a first storage device and a second volume on a second storage device may also include executable code that returns a result indicating successfully reversing the communication path.

[0013] According to the present invention, determining device characteristics includes obtaining a first globally accessible value, if the first globally accessible value corresponds to a stored first value, obtaining device characteristics data from a relatively fast memory, if the first globally accessible value does not correspond to the stored first value, obtaining a second globally accessible value, if the second globally accessible value corresponds to a stored second value, obtaining device characteristics data from a relatively fast memory, if the second globally accessible value does not correspond to the stored second value, obtaining device characteristics data from a relatively slow memory and updating the relatively fast memory, the stored first value, and the stored second value. The globally accessible first value may include device I/O information. The globally acces-

sible values may be stored in global memory that is accessible to a plurality of processors. The relatively slow memory may be global memory that is accessible to a plurality of processors. The relatively fast memory may be local to a processor that accesses the first and second stored values. The second stored value may indicate whether a device is a source or a destination for data in a remote data communication link.

[0014] According further to the present invention, accessing a device includes obtaining status information that indicates how to access the device, examining at least a portion of the status information to determine if configuration of the device has changed, if the configuration of the device has not changed, accessing stored configuration data from relatively fast memory, and if the configuration of the device has changed, accessing new configuration information from relatively slow memory and updating the stored configuration data with the new configuration information. The relatively slow memory may be global memory that is accessible to a plurality of processors. The relatively fast memory may be local to a processor that accesses the fast memory. The stored configuration data may indicate whether a device is a source or a destination for data in a remote data communication link. Examining at least a portion of the status information may include determining whether the portion of the status information is equal to a first stored value. Accessing a device may also include, if the portion of the status information is not equal to the first stored value, obtaining additional status information and determining if a portion of the additional status information is equal to a second stored value. Accessing a device may also include following updating the stored configuration data, updating the first and second stored values.

[0015] According further to the present invention, a computer program product that determines device characteristics, includes executable code that obtains a first globally accessible value, executable code that obtains device characteristics data from a relatively fast memory if the first globally accessible value corresponds to a stored first value, executable code that obtains a second globally accessible value if the first globally accessible value does not correspond to the stored first value, executable code that obtains device characteristics data from a relatively fast memory if the second globally accessible value corresponds to a stored second value, and executable code that obtains device characteristics data from a relatively slow memory and updates the relatively fast memory, the stored first value, and the stored second value if the second globally accessible value does not correspond to the stored second value. The globally accessible first value may include device I/O information. The second stored value may indicate whether a device is a source or a destination for data in a remote data communication link.

[0016] According further to the present invention, a computer program product that accesses a device in-

cludes executable code that obtains status information that indicates how to access the device, executable code that examines at least a portion of the status information to determine if configuration of the device has changed, executable code that accesses stored configuration data from relatively fast memory if the configuration of the device has not changed, and executable code that accesses new configuration information from relatively slow memory and updates the stored configuration data with the new configuration information if the configuration of the device has changed. The stored configuration data may indicate whether a device is a source or a destination for data in a remote data communication link. Executable code that examines at least a portion of the status information may include executable code that determines whether the portion of the status information is equal to a first stored value. A computer program product that accesses a device may also include executable code that obtains additional status information and determines if a portion of the additional status information is equal to a second stored value if the portion of the status information is not equal to the first stored value. A computer program product that accesses a device may also include executable code that updates the first and second stored values following updating the stored configuration data.

[0017] According further to the present invention, an apparatus that determines device characteristics includes means for obtaining a first globally accessible value, means for obtaining device characteristics data from a relatively fast memory if the first globally accessible value corresponds to a stored first value, means for obtaining a second globally accessible value if the first globally accessible value does not correspond to the stored first value, means for obtaining device characteristics data from a relatively fast memory if the second globally accessible value corresponds to a stored second value, means for obtaining device characteristics data from a relatively slow memory and updating the relatively fast memory, the stored first value, and the stored second value if the second globally accessible value does not correspond to the stored second value. The globally accessible first value may include device I/O information. The globally accessible values may be stored in global memory that is accessible to a plurality of processors. The relatively slow memory may be global memory that is accessible to a plurality of processors. The relatively fast memory may be local to a processor that accesses the first and second stored values. The second stored value may indicate whether a device is a source or a destination for data in a remote data communication link.

[0018] According further to the present invention an apparatus that accesses a device, includes means for obtaining status information that indicates how to access the device, means for examining at least a portion of the status information to determine if configuration of the device has changed, means for accessing stored

configuration data from relatively fast memory if the configuration of the device has not changed, and means for accessing new configuration information from relatively slow memory and updating the stored configuration data with the new configuration information if the configuration of the device has changed. The relatively slow memory may be global memory that is accessible to a plurality of processors. The relatively fast memory may be local to a processor that accesses the fast memory. The stored configuration data may indicate whether a device is a source or a destination for data in a remote data communication link. Means for examining at least a portion of the status information may include means for determining whether the portion of the status information is equal to a first stored value. The apparatus may further include means for obtaining additional status information and determining if a portion of the additional status information is equal to a second stored value if the portion of the status information is not equal to the first stored value. The apparatus may further include means for updating the first and second stored values following updating the stored configuration data.

Brief Description of Drawings

[0019] Figure 1 is a schematic diagram showing a host, a local storage device, and a remote data storage device used in connection with the system described herein.

[0020] Figure 2 is a flow chart illustrating operation of the system described herein.

[0021] Figure 3 is a flow chart illustrating determining which volumes to use for read and write operations according to the system described herein.

[0022] Figure 4 is a diagram illustrating global memory for storage devices according to the system described herein.

[0023] Figure 5 is a flow chart illustrating a hierarchy for determining whether to use static or dynamic configuration data in connection with assessing RDF configuration of a device according to the system described herein.

[0024] Figure 6 is a schematic diagram showing in detail a storage device and components thereof used in connection with the system described herein.

[0025] Figure 7 is a diagram illustrating a ready buffer and a set of bytes used in connection with the system described herein.

[0026] Figure 8 is a flow chart illustrating steps performed in connection with accessing the ready buffer and the set of bytes set forth in Figure 7.

[0027] Figure 9 is a flow chart illustrating an alternative embodiment of the system described herein.

Detailed Description of Various Embodiments

[0028] Referring to Figure 1, a diagram 20 shows a relationship between a host 22, a local storage device

24 and a remote storage device 26. The host 22 reads and writes data from and to the local storage device 24 via a host adapter 28, which facilitates the interface between the host 22 and the local storage device 24. Data from the local storage device 24 is copied to the remote storage device 26 via an RDF link 29 to cause the data on the remote storage device 26 to be identical to the data on the local storage device 24. Note that there may be a time delay between the transfer of data from the local storage device 24 to the remote storage device 26, so that the remote storage device 26 may, at certain points in time, contain data that is not identical to the data on the local storage device 24. Communication between data storage devices or systems, such as Symmetrix™ data storage systems by EMC Corporation, may be facilitated, for example, using the RDF product, as also provided by EMC Corporation.

[0029] The local storage device 24 includes an RDF adapter unit (RA) 30 and the remote storage device 26 includes an RA 32. The RA's 30, 32 are coupled to the RDF link 29 and are similar to the host adapter 28, but are used to transfer data between the storage devices 24, 26. The software used in connection with the RA's 30, 32 is discussed in more detail hereinafter.

[0030] The storage devices 24, 26 may include one or more disks, each containing a different portion of data stored on each of the storage devices 24, 26. Figure 1 shows the storage device 24 including a plurality of disks 33a, 33b, 33c and the storage device 26 including a plurality of disks 34a, 34b, 34c. The RDF functionality described herein may be applied so that the data for at least a portion of the disks 33a-33c of the local storage device 24 is copied, using RDF, to at least a portion of the disks 34a-34c of the remote storage device 26. It is possible that other data of the storage devices 24, 26 is not copied between the storage devices 24, 26, and thus is not identical.

[0031] Each of the disks 33a-33c is coupled to a corresponding disk adapter unit (DA) 35a, 35b, 35c that provides data to a corresponding one of the disks 33a-33c and receives data from a corresponding one of the disks 33a-33c. Similarly, a plurality of DA's 36a, 36b, 36c of the remote storage device 26 are used to provide data to corresponding ones of the disks 34a-34c and receive data from corresponding ones of the disks 34a-34c. A data path exists between the DA's 35a-35c, the HA 28 and RA 30 of the local storage device 24. Similarly, a data path exists between the DA's 36a-36c and the RA 32 of the remote storage device 26.

[0032] The local storage device 24 also includes a global memory 37 that may be used to facilitate data transferred between the DA's 35a-35c, the HA 28 and the RA 30. The memory 37 may contain parameters from system calls, tasks that are to be performed by one or more of the DA's 35a-35c, the HA 28 and the RA 30, and a cache for data fetched from one or more of the disks 33a-33c. Similarly, the remote storage device 26 includes a global memory 38 that may contain parame-

ters from system calls, tasks that are to be performed by one or more of the DA's 36a-36c and the RA 32, and a cache for data fetched from one or more of the disks 34a-34c. Use of the memories 37, 38 is described in more detail hereinafter.

[0033] The storage space in the local storage device 24 that corresponds to the disks 33a-33c may be subdivided into a plurality of volumes or logical devices. The logical devices may or may not correspond to the physical storage space of the disks 33a-33c. Thus, for example, the disk 33a may contain a plurality of logical devices or, alternatively, a single logical device could span both of the disks 33a, 33b. Similarly, the storage space for the remote storage device 26 that comprises the disks 34a-34c may be subdivided into a plurality of volumes or logical devices, where each of the logical devices may or may not correspond to one or more of the disks 34a-34c.

[0034] Providing an RDF mapping between portions of the local storage device 24 and the remote storage device 26 involves setting up a logical device on the remote storage device 26 that is a remote mirror for a logical device on the local storage device 24. The host 22 reads and writes data from and to the logical device on the local storage device 24 and the RDF mapping causes modified data to be transferred from the local storage device 24 to the remote storage device 26 using the RAs, 30, 32 and the RDF link 29. In steady state operation, the logical device on the remote storage device 26 contains data that is identical to the data of the logical device on the local storage device 24. The logical device on the local storage device 24 that is accessed by the host 22 is referred to as the "R1 volume" (or just "R1") while the logical device on the remote storage device 26 that contains a copy of the data on the R1 volume is called the "R2 volume" (or just "R2"). Thus, the host reads and writes data from and to the R1 volume and RDF handles automatic copying and updating of the data from the R1 volume to the R2 volume.

[0035] In some situations, it may be beneficial to allow the host 22 to create and destroy RDF volumes during operation of the system. Note that RDF volumes may be created and destroyed in pairs so that an R1/R2 pair may be destroyed or an R1/R2 pair may be created. Creating or destroying R1/R2 pairs may be initiated by the host 22. The host may send a multihop/multiexecute system command to data storage devices. The multihop/multiexecute system command is a single system command that is provided to multiple storage devices and indicates operations to be performed by the multiple storage devices. The command allows for a host computer to issue a remote system call indirectly to one or more data storage devices through layers of other data storage devices. The command may indicate that any one or more of the data storage devices is to perform the command. For example, the host 22 may send a multihop/multiexecute system command requesting that a particular R1/R2 pair be destroyed where the R1

volume is on the local storage device 24 and the R2 volume is on the remote storage device 26 by having each of the storage devices 24, 26 locally modify a table (discussed in more detail below) that is used internally by each of the storage devices 24, 26 to govern setting up and managing RDF volumes. Creating an R1/R2 pair involves creating the R1 volume on one storage device and creating the R2 volume on another storage device.

[0036] Further note that, in certain instances, it may be useful to establish an RDF connection to one or more existing volumes. Thus, for example, an existing source volume may be connected to newly-created destination volume using RDF. Accordingly, for the discussion set forth herein, references to creating volumes may be understood to include creating connections to existing volumes. Similarly, references to destroying volumes may be understood to include simply destroying RDF connections thereto, as appropriate. Note that, both in the case of creating or destroying new volumes and in the case of creating or destroying simply the connections to existing volume, the same tables (discussed below) that contain RDF connection data are modified.

[0037] Referring to Figure 2, a flow chart 50 illustrates steps performed in connection with creating or destroying R1/R2 pairs. Processing begins at a first step 52 where it is determined if a destroy command is being issued. In some embodiments, only create or destroy commands are issued. Thus, if a command is not a destroy command, the command is a create command. If it is determined at the step 52 that a destroy command was issued, then control passes from the step 52 to a step 54 any where background I/O operations, such as background copying, are suspended. Prior to destroying the R1/R2 pair, it is useful to first suspend any background I/O. In other embodiments, the step 54 is not necessary and not executed because destroying the volumes will cause background I/O (and other operations involving the R1/R2 pair) to cease automatically.

[0038] Following the step 54 is a step 56 where an allocation table on the storage device corresponding to the R1 volume is modified. The allocation table contains dynamic information about the RDF configuration of a storage device. The allocation table may contain a two dimensional array indexed by logical device identifiers (such as numbers) and by mirror numbers for each of the logical devices. In some embodiments, each device may have up to four mirrors. Other embodiments may employ more than four mirrors. Entries for the allocation table may indicate whether a mirror for a device is a local mirror, an R1 volume, an R2 volume, or some other type of mirror, such as BCV or a RAID mirror.

[0039] At the step 56, the allocation table on the storage device containing the R1 volume is modified to remove the R1 volume. Following the step 56 is a step 58 where the allocation table on the storage device containing the R2 volume is modified to remove the R2 volume. Following the step 58 is a step 60 where the result of the previously executed operations (e.g., success or

failure) is returned. Following the step 60, processing is complete.

[0040] If it is determined at the step 52 that a destroy command is not being issued (and hence a create command is being issued), then control passes from the step 52 to a step 62 where the R1 or R2 volume is created at a first site. In one embodiment, the host issues the multihop/multiexecute command to a first storage device such as the local storage device 24, in which case the first site would be the local storage device 24 (i.e., the first site to receive the command). In an embodiment illustrated herein, the first site on which the creation of an R1/R2 pair is attempted is the local storage device 24 that is coupled directly to the host 22. Creation at the first site includes modifying the appropriate allocation table. Following the step 62 is a test step 64 where it is determined if creation of the R1 or R2 volume at the step 62 was successful. If not, control passes from the step 64 to a step 66 where an error indication is returned to the host 22. Following the step 66, processing is complete.

[0041] If it is determined at the test step 64 that the creation of the R1 or R2 volume at the step 62 was successful, then control passes from the step 64 to a step 68 where the R1 or R2 volume, as the case may be, is created at a second site. Creation at the second site includes modifying the appropriate allocation table. The second site may be the second storage device to receive the command. In an embodiment illustrated herein, the second site is the remote storage device 26 that receives a command from the host 22 through the local storage device 24.

[0042] Following the step 68 is a test step 70 where it is determined if creation at the step 68 was successful. If it is determined at the test step 70 that the creation at the second site was not successful, then control passes from the step 70 to a step 72 where the first site that was created is destroyed. Note that R1 and R2 volumes are created in pairs so that failure to create one of the volumes at the second site causes the volume that was successfully created at the first site to be destroyed at the step 72. Following the step 72 is a step 74 where an error is returned to the host 22 in a manner similar to returning an error at the step 66, discussed above. Following the step 74, processing is complete.

[0043] If it is determined at the test step 70 that creation of the R1 or R2 volume at the second site at the step 68 was successful, control transfers from the step 70 to a step 75 where invalid tracks for each of the storage devices 24, 26 are set. The invalid tracks are used to indicate differences between the R1 and R2 volumes. In instances where the R1 volume contains the initial data, all of the tracks of the R1 volume are set to invalid indicating that all of the R1 data is inconsistent with the R2 data, and thus needs to be copied from R1 to R2. If the R2 volume contains the initial data to be used, then all of the tracks of the R2 volume are invalidated and a Sync command is provided to the storage device con-

taining the R1 volume to cause the tracks of the local mirror corresponding to the R1 volume to be invalidated, thus initiating a background copy of data from R2 to R1. Following the step 75 is a step 78 where a not ready indicator for the R1 volume is cleared, indicating that RDF operations between R1 and R2 may commence. Note that the not ready indicator refers to the R1/R2 link, and not necessarily to the ability of R1 and/or R2 to accept data from a host. Note that data can be organized in units other than tracks, in which case, the operations disclosed herein would be performed on whatever units other than tracks that are used.

[0044] Following the step 78 is a step 79 where an indicator that the operation was successful is returned to the host 22. Following the step 79, processing is complete. Note that setting invalid tracks and setting and clearing device not ready status, discussed above, may be performed in a different order than that indicated in Figure 2 and discussed above, provided that the settings are maintained in an appropriate state during the creation operation and thereafter.

[0045] In some embodiments, it may be useful to create and destroy the R1/R2 volumes in a particular order. For example, it may be useful to create the R2 volume first followed by the R1 volume. Similarly, it may be useful to destroy the R1 volume first followed by destroying the R2 volume. Alternatively, the volumes may be created and destroyed in any order while operations may be controlled by setting a not ready indicator for the R1 volume, where the not ready indicator for the R1 volume controls whether operations take place. Thus, the R1 and R2 volumes may be created in any order, but the not ready indicator of the R1 volume is not set to indicate ready unless and until both R1 and R2 have been created. Similarly, setting the not ready indicator for R1 as a first step allows subsequent steps for destroying R1 and R2 to take place in any order.

[0046] In some embodiments, it may be useful to have the command that creates the R1/R2 pair indicate the specific volumes to be used for both the R1 volume and the R2 volume. That is, the host 22 may select an existing logical device from the local storage device 24 as the R1 volume and, in addition, may select an existing logical device from the remote storage device 26 as an R2 volume. In an alternative embodiment, it may be possible to have the remote storage device 26 select any unused logical device as the R2 volume. Alternatively still, an API may be layered on to one of the storage devices 24, 26 and/or the host 22 to pick the R2 volume in instances where the system command expects specific volume identification for the R2 volume.

[0047] Note, in addition, there may be situations where the initial data for a newly-created R1/R2 pair is found on the R2 volume. For example, if the remote storage device 26 has coupled thereto a host (not shown) that fails, it may be useful to restart the system with the host 22 coupled to the local storage device 24. However, in that case, the initial data for starting up may be located

on the R2 volume (i.e., on the remote storage device 26). Thus, at start up, the initial data is copied from the R2 device to the R1 device, after which the host 22 maintains normal RDF operation with the R1 device being located on the local storage device 24 and the R2 device being located on the remote storage device 26. The copy from the R2 device to the R1 device (or, from an R1 device to an R2 device) may be provided using a background copy.

[0048] Referring to Figure 3, a flow chart 80 illustrates operation of the system after an R1/R2 pair is initialized but prior to completion of the background copy. Note that the host 22 may read and write from and to an R1 volume even if the background copy has not completed.

[0049] Processing begins at an initial test step 82 where it is determined if the R2 volume contains the initial data. If not, then control passes from the test step 82 to a step 84 where the R1 volume is used for the read or write operation requested by the host 22.

[0050] If it is determined at the test step 82 that the R2 volume is specified as containing the initial data, then control passes from the test step 82 to a test step 86 where it is determined if the track being written to or read from is invalid. An invalid track indicates that the initial data, located on the R2 volume, has not yet been copied back to the R1 volume. Setting invalid tracks is discussed above in connection with Figure 1. Thus, if it is determined at the test step 86 that the track being written to or read from is not invalid, control passes from the step 86 to a step 88 where the R2 volume is used for the read or write operation using RDF. Note that, in some instances, less than an entire track is written by the host. Thus, for an invalid track where R2 contains the initial data, it is useful to have the data from the host 22 provided to the R2 volume. Following the step 88, processing is complete. Alternatively, if it is determined at the test step 86 that the track being read to or written from is not invalid, then control passes from the test step 86 to the step 84 where the R1 volume is used, in a conventional fashion. Following the step 84, processing is complete.

[0051] Referring to Figure 4, a diagram shows global memory that could represent the memory 37 of the local storage device 24 or the memory 38 of the remote storage device 26. The global memory includes a memory location for static configuration data 92 and a memory location for dynamic configuration data 94, such as the allocation table, discussed above. Static configuration data 92 includes configuration information for the storage device that is set up at the factory or by a technician. The static configuration data 92 may be provided in a file that is read into global semiconductor memory or it may be provided in non-volatile portions of semiconductor memory.

[0052] The dynamic configuration data 94 represents the configuration of the corresponding system that has been modified from the static configuration data 92. Thus, for example, the static configuration data 92 may

indicate that particular volumes form an R1/R2 RDF pair while the dynamic configuration data 94 overrides the static configuration data 92 by indicating that the particular R1/R2 RDF pair has been subsequently destroyed.

The dynamic configuration data 94 may also be stored on a disk and read into electronic memory and/or may be stored in electronic memory that is non-volatile. In some embodiments, the dynamic configuration data 94 may override the static configuration data 92. In other embodiments, the dynamic configuration data 94 is only additive and may only be used in instances where there is no entry in the static configuration data 92 corresponding to an item.

[0053] Referring to Figure 5, a flow chart 100 illustrates determination of a configuration of a particular logical device. Note that each read or write access from or to a logical device uses knowledge of the configuration of the device to determine if, for example, the device is part of an R1/R2 pair.

[0054] Processing begins at a first step 102 where it is determined if there is an entry for the logical device in the dynamic configuration data 94. If not, then control passes from the step 102 to a step 104 where the static configuration data 92 is consulted to determine the configuration of the device. Note that it is possible for each logical device to have a local, and thus easily accessible, copy of the static configuration data 92 since, by definition, the static configuration data 92 does not change. Following the step 104, processing is complete.

[0055] If it is determined at the test step 102 that there is an entry for the logical device in the dynamic configuration data 94, then control passes from the test step 102 to a step 106 where the dynamic configuration data 94 is used to determine the configuration of the logical device. Note that, as an alternative, it is possible to first check the static configuration data 92 and then check to see if an overriding entry exists in the dynamic configuration data 94.

[0056] Referring to Figure 6, a diagram shows a storage device 110 in more detail. The storage device 110 includes a plurality of host adapters 112-114 and a plurality of disk adapters 116-118. Each of the disk adapters 116-118 is coupled to a respective one of a plurality of disks 120-122. The storage device 110 also includes a global memory 124 an RDF adapter (RA) 126 and an outside connection 128 to the RA 126. A bus 130 connects the HAs 112-114, the DAs 116-118, the global memory 124 and the RA 126. Each of the HAs 112-114 includes a connection for coupling to a host (not shown).

[0057] For every read and write operation performed by one of the hosts coupled to one of the HAs 112-114, the corresponding one of the HAs 112-114 determines where the data is to be placed. For example, if the storage device 110 includes a local volume of an R1/R2 pair, the particular one of the HAs 112-114 that receives the data from the host must be able to direct the data to the correct one of the disk adapters 116-118 in order to provide the data to the R1 volume and must also be able

to appropriately direct the data to the RA 126 that provides the data to the R2 volume on another storage device (not shown). In addition, in connection with a background copy operation, the DAs 116-118 access information indicating the source and destination of the data. Thus, for example, if an R1/R2 pair is set up between the storage device 110 and another remote storage device (not shown), then the DAs 116-118 would assist in the background copy of data from the volume containing the initial data (R1 or R2) to the other volume. In such a case, the DAs 116-118 access information that indicates where the data should go (i.e., which volume is the corresponding one of the volumes of the R1/R2 pair). Similarly, it is useful for the RA 126 to access the configuration information for the volumes in order to be able to process incoming data, for example.

[0058] Accordingly, it is useful for the HAs 112-114, the DAs 116-118 and the RA 126 to be able to have access to accurate information indicating the set up configuration of the volumes on the storage device 110. One way to obtain the information is to read the dynamic configuration data 94 from global memory 124 for each read or write operation and then, if the device for which the inquiry is made is not found in the dynamic configuration data 94, to then access the static configuration data 92. However, accessing global memory for every I/O operation could adversely impact the performance of the storage device 110. Accordingly, a more efficient mechanism for determining the configuration of the devices of the storage device 110 is provided.

[0059] Referring to Figure 7, a diagram 150 illustrates a mechanism for decreasing the number of accesses to global memory used to determine configuration data. A ready buffer 152 represents data that is accessed each time an I/O is performed. The ready buffer 152 indicates useful I/O related information, such as the state of the devices and/or the state of the system. Thus, the ready buffer 152 is accessed for each I/O operation.

[0060] The ready buffer 152 include a revision number field 154 which, in some embodiments, is a byte of data that is incremented each time a device configuration has changed. Thus, a device performing an I/O can look to the revision number field 154 to determine if further inquiry is necessary.

[0061] A set of bytes 156 may be accessed in response to the revision number field 154 changing. A particular byte 158 of the set of bytes 156 could be used to indicate a change in the dynamic RDF configuration for the storage device. Thus, the combination of the ready buffer 152 and the set of bytes 156 may be used to minimize accesses to global memory 124 each time an I/O is performed.

[0062] Referring to Figure 8, a flow chart 170 illustrates steps performed in connection with using the ready buffer 152 and the set of bytes 156. Processing begins at a first step 172 where a device performing an I/O determines if the ready buffer field 154 has changed (e.g., incremented). If it is determined at the step 172

that the ready buffer field has not changed, then control passes from the step 172 to a step 174 where the device, such as one of the HAs 112-114, one of the DAs 116-118, and/or the RA 126, uses a locally stored data to determine configuration information. That is, the revision number field 154 being equal to a previously observed version of the revision number field 154 indicates that no changes have been made to the RDF configuration. Thus, control passes from the step 172 to the step 174 where a locally stored copy of the configuration data may be used by the device performing the I/O in order to determine how to handle particular data. Following the step 174, processing is complete.

[0063] If it is determined at the test step 172 that the revision number field 154 has changed, then control passes from the step 172 to a step 176 where the set of bytes 156 is fetched. The set of bytes 156 may be stored, for example, in global memory 124 or in another memory, such as non-volatile electronic memory or on a disk.

[0064] Following the step 176 is a test step 178 where it is determined if the byte 158 in the set of bytes 156 corresponding to an RDF configuration change has been modified (e.g., incremented). If not, then control passes from the step 178 to the step 174, discussed above, where the stored local copy of the configuration data is used to determine the RDF configuration. Following step 174, processing is complete.

[0065] If it is determined at the test step 178 that the RDF configuration has changed, then control passes from the step 178 to a step 180 where the dynamic configuration data 94 in the global memory 124 is accessed. Following the step 180 is a step 182 where the stored local copy of the configuration data is updated with the new version of the data from global memory 124. Following the step 182 is the step 174 where the locally stored version of the configuration data (which has just been updated) is used by the device to determine the RDF configuration. Following the step 174, processing is complete.

[0066] In some embodiments, it is possible to also include a background mechanism that periodically accesses the dynamic configuration data from the allocation table and updates the local copy of the configuration data in the storage for each device. Thus, each device may be updated by this background task without first having to perform an I/O.

[0067] Note that the mechanism discussed herein for using revision numbers to control accesses to global memory when determining the RDF configuration of the system may be applied to other system parameters as well, such as BCV configuration. A BCV (Business Continuity Volume) may contain a consistent, synchronized copy of data from another device configured for use, for example, with host applications. The BCV copy of the data may be used independently by another application to act on the data copy while letting the host perform continued data operations to the other device.

[0068] Referring to Figure 9, a flow chart 50' illustrates an alternative embodiment to the embodiment described in connection with the flow chart 50 of Figure 2. The flow chart 50 of Figure 2 shows a system where there are two commands, create and destroy. The system illustrated by the flow chart 50' of Figure 9 provides for a third command that swaps R1 and R2 in addition to the create and destroy commands discussed above. Swapping R1 and R2 may be performed using a single command that causes the R1 volume to become the R2 volume and causes the R2 volume to become the R1 volume. This may be useful in instances where, for example, a host coupled to a local storage device is replaced by a different host that is coupled to a corresponding remote storage device. The command may be provided by the host. The host may send a multihop/multiexecute system command, as also described elsewhere herein. Processing begins at a first step 52' where it is determined if a destroy command has been issued. If so, then control passes from the step 52' to the step 54, which is discussed above in connection with the flow chart 50 of Figure 2. If, on the other hand, it is determined at the step 52' that a destroy command has not been issued, then control passes from the step 52' to a step 53 where it is determined if a create command has been issued. If a create command has been issued, control passes from the step 53 to the step 62, which is discussed above in connection with the flow chart 50 of Figure 2.

[0069] If it is determined at the test step 53 that a create command has not been issued, then a swap R1/R2 command has been issued. That is, in a system having three commands, a process of elimination provides that if it is determined at the step 52' that a destroy command has not been issued, and if it is determined at the step 53 that a create command has not been issued, then a swap R1/R2 command has been issued. Alternatively, a particular portion of code may be executed based on the command received, which could be a swap R1/R2 command. For the swap R1/R2 command, control transfers from the step 53 to a step 192 where the device not ready flag for R1 is set. Setting the device not ready flag for R1 at the step 192 causes RDF operations to cease for the R1/R2 pair. Following the step 192 is a step 194 where an entry for the allocation table corresponding to the storage device containing the R1 volume and an entry for the allocation table corresponding to the storage device containing the R2 volume are both modified so that the R1 volume becomes the R2 volume and vice versa. Following the step 194 is a step 196 where the device not ready bit is cleared for the new R1 device. Clearing the device not ready bit for the R1 volume at the step 196 allows RDF operations to resume. However, when RDF operations resume, the previous R2 volume is now the R1 volume and the previous R1 volume is now the R2 volume. Following the step 196 is a step 198 where the result of the operations corresponding to the command for swapping R1 and R2 (e.g., success or

failure) are returned to the host. Following the step 198, processing is complete.

[0070] Note that for the system disclosed herein, the host 22 may be a standalone computing device, part of a network, another storage device, a computer having some storage capabilities and/or any device capable of providing the functionality disclosed herein. Similarly, the system disclosed herein may be practiced with any appropriate device used for the local storage device 24, including possibly another computing device, network connection, etc. configured to emulate the local storage device 24 and/or otherwise provide the functionality disclosed herein.

[0071] While the invention has been disclosed in connection with various embodiments, modifications thereon will be readily apparent to those skilled in the art. Accordingly, the spirit and scope of the invention is set forth in the following claims.

Claims

1. A method of determining device characteristics, comprising:

obtaining a globally accessible value;

further **characterized by** comprising:

obtaining a first globally accessible value (154);
if the first globally accessible value (154) corresponds to a stored first value, obtaining device characteristics data from a relatively fast memory;
if the first globally accessible value (154) does not correspond to the stored first value, obtaining a second globally accessible value (158);
if the second globally accessible value (158) corresponds to a stored second value, obtaining device characteristics data from a relatively fast memory;
if the second globally accessible value (158) does not correspond to the stored second value, obtaining device characteristics data from a relatively slow memory (124) and updating the relatively fast memory, the stored first value, and the stored second value.

2. A method, according to claim 1, wherein the globally accessible first value includes device I/O information.
3. A method, according to claim 1 or 2, wherein the globally accessible values are stored in global memory that is accessible to a plurality of processors.
4. A method, according to any one of the preceding

claims, wherein the relatively slow memory is global memory that is accessible to a plurality of processors.

5. A method, according to any one of the preceding claims, wherein the relatively fast memory is local to a processor that accesses the first and second stored values.

6. A method, according to any one of the preceding claims, wherein the second stored value indicates whether a device is a source or a destination for data in a remote data communication link.

7. A method of accessing a device, comprising:
obtaining information;

further **characterized by** comprising:

obtaining status information that indicates how to access the device;
examining at least a portion of the status information to determine if configuration of the device has changed;
if the configuration of the device has not changed, accessing stored configuration data from relatively fast memory; and
if the configuration of the device has changed, accessing new configuration information from relatively slow memory and updating the stored configuration data with the new configuration information.

8. A method, according to claim 7, wherein the relatively slow memory is global memory that is accessible to a plurality of processors.

9. A method, according to claim 7 or 8, wherein the relatively fast memory is local to a processor that accesses the fast memory.

10. A method, according to any one of claims 7 to 9, wherein the stored configuration data indicates whether a device is a source or a destination for data in a remote data communication link.

11. A method, according to any one of claims 7 to 10, wherein examining at least a portion of the status information includes determining whether the portion of the status information is equal to a first stored value.

12. A method, according to claim 11, further comprising:
if the portion of the status information is not equal to the first stored value, obtaining additional status information and determining if a

portion of the additional status information is equal to a second stored value.

13. A method, according to claim 12, further comprising:

following updating the stored configuration data, updating the first and second stored values.

14. A computer program product that determines device characteristics, comprising executable code for executing the method according to any one of the preceding claims.

15. Apparatus that determines device characteristics, comprising:

means for obtaining a globally accessible value;

further **characterized by** comprising:

means for obtaining a first globally accessible value (154);
means for obtaining device characteristics data from a relatively fast memory if the first globally accessible value (154) corresponds to a stored first value;
means for obtaining a second globally accessible value (158) if the first globally accessible (154) value does not correspond to the stored first value;
means for obtaining device characteristics data from a relatively fast memory if the second globally accessible value (158) corresponds to a stored second value;
means for obtaining device characteristics data from a relatively slow memory (124) and updating the relatively fast memory, the stored first value, and the stored second value if the second globally accessible value (158) does not correspond to the stored second value.

16. Apparatus, according to claim 15, wherein the globally accessible first value includes device I/O information.

17. Apparatus, according to claim 15 or 16, wherein the globally accessible values are stored in global memory that is accessible to a plurality of processors.

18. Apparatus, according to any one of claims 14 to 16, wherein the relatively slow memory is global memory that is accessible to a plurality of processors.

19. Apparatus, according to any one of claims 14 to 17, wherein the relatively fast memory is local to a processor.

essor that accesses the first and second stored values.

means for updating the first and second stored values following updating the stored configuration data.

20. Apparatus, according to any one of claims 14 to 18, wherein the second stored value indicates whether a device is a source or a destination for data in a remote data communication link. 5
21. Apparatus that accesses a device, comprising: 10
- means for obtaining information;
- further **chracterized by** comprising:
- means for obtaining status information that indicates how to access the device; 15
- means for examining at least a portion of the status information to determine if configuration of the device has changed;
- means for accessing stored configuration data from relatively fast memory if the configuration of the device has not changed; and 20
- means for accessing new configuration information from relatively slow memory and updating the stored configuration data with the new configuration information if the configuration of the device has changed. 25
22. Apparatus, according to claim 21, wherein the relatively slow memory is global memory that is accessible to a plurality of processors. 30
23. Apparatus, according to claim 21 or 22, wherein the relatively fast memory is local to a processor that accesses the fast memory. 35
24. Apparatus, according to any one of claims 21 to 23, wherein the stored configuration data indicates whether a device is a source or a destination for data in a remote data communication link. 40
25. Apparatus, according to any one of claims 21 to 26, wherein means for examining at least a portion of the status information includes means for determining whether the portion of the status information is equal to a first stored value. 45
26. Apparatus, according to claim 25, further comprising: 50
- means for obtaining additional status information and determining if a portion of the additional status information is equal to a second stored value if the portion of the status information is not equal to the first stored value. 55
27. Apparatus, according to claim 26, further comprising:

20 →

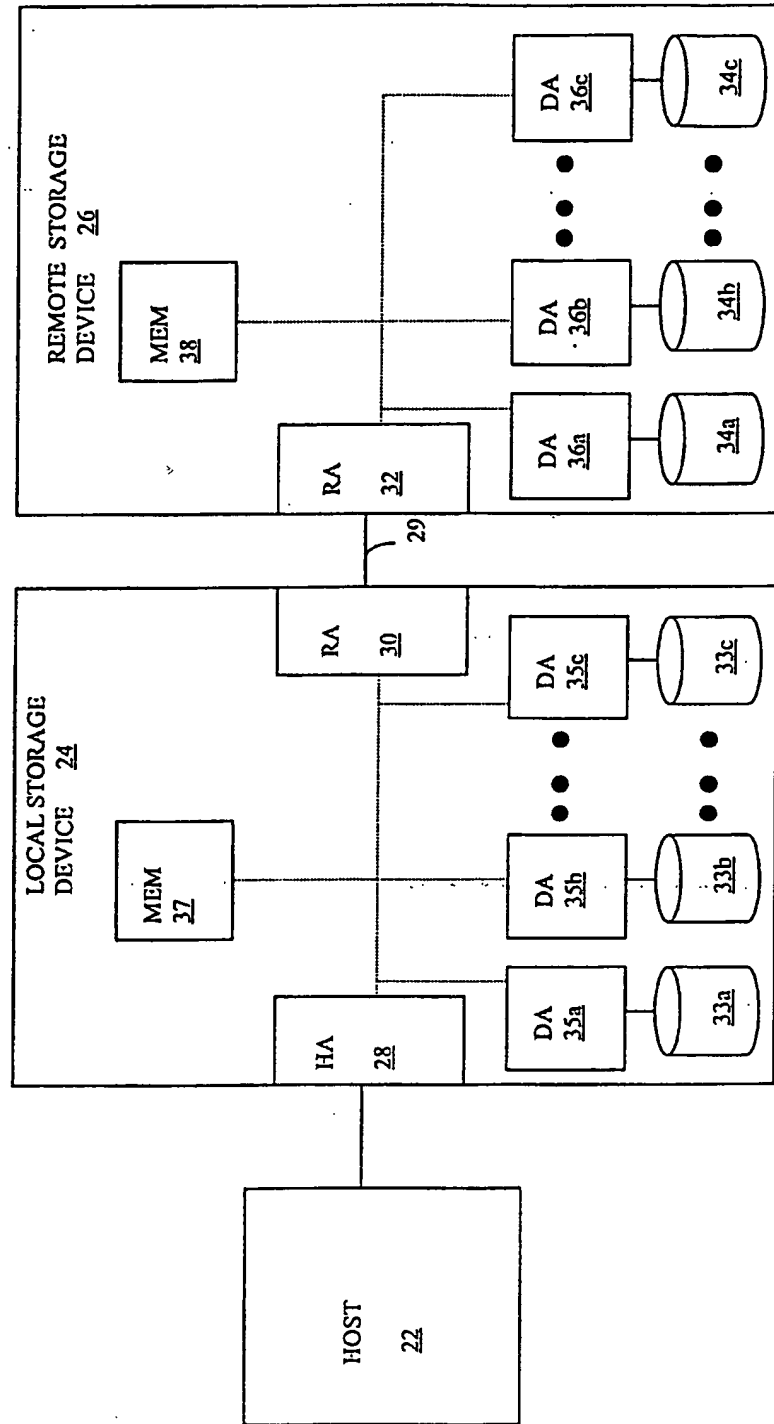


FIGURE 1

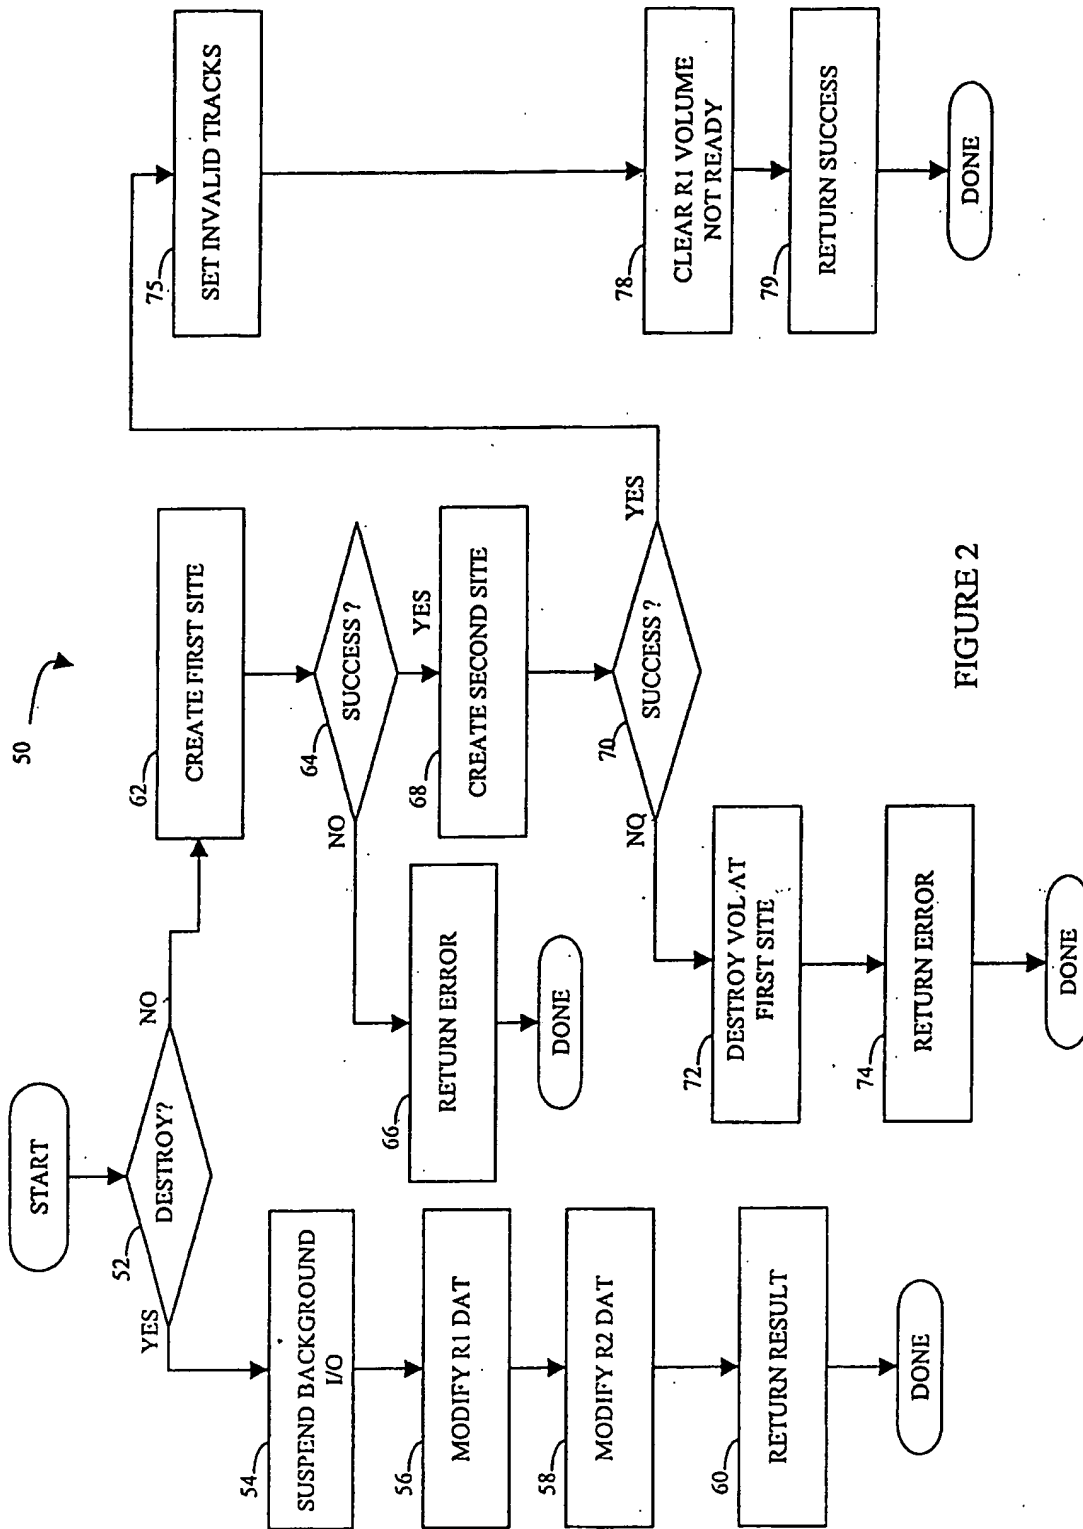


FIGURE 2

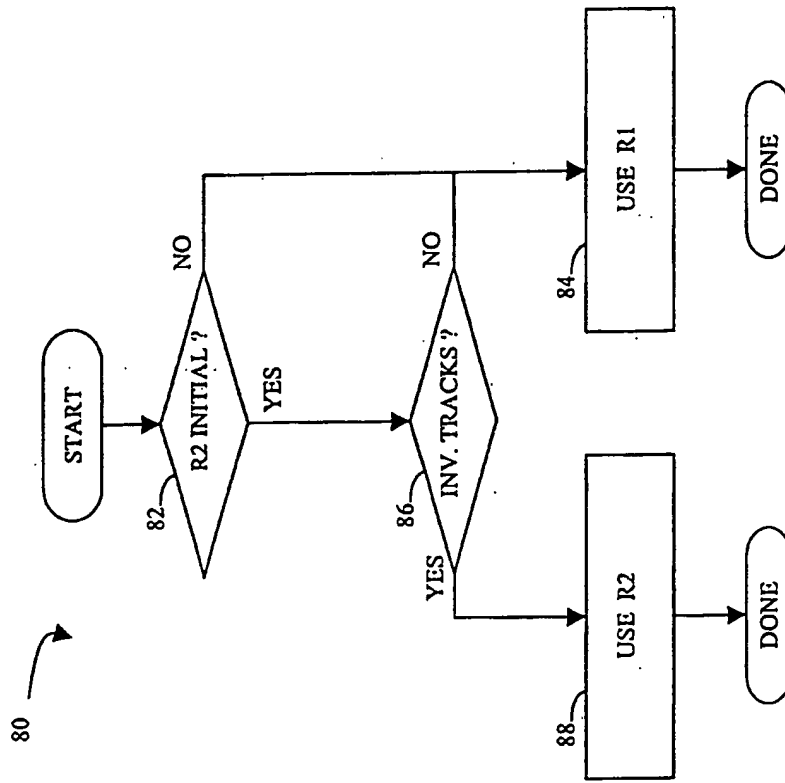


FIGURE 3

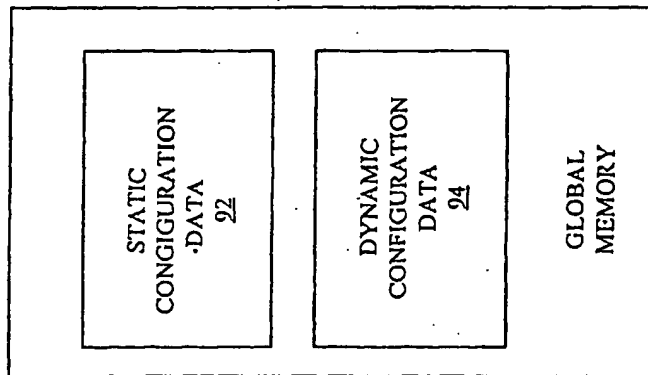


FIGURE 4

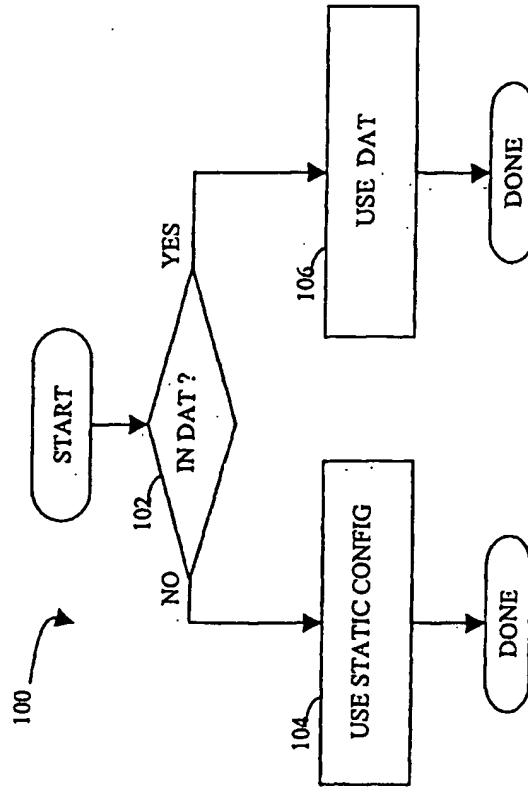


FIGURE 5

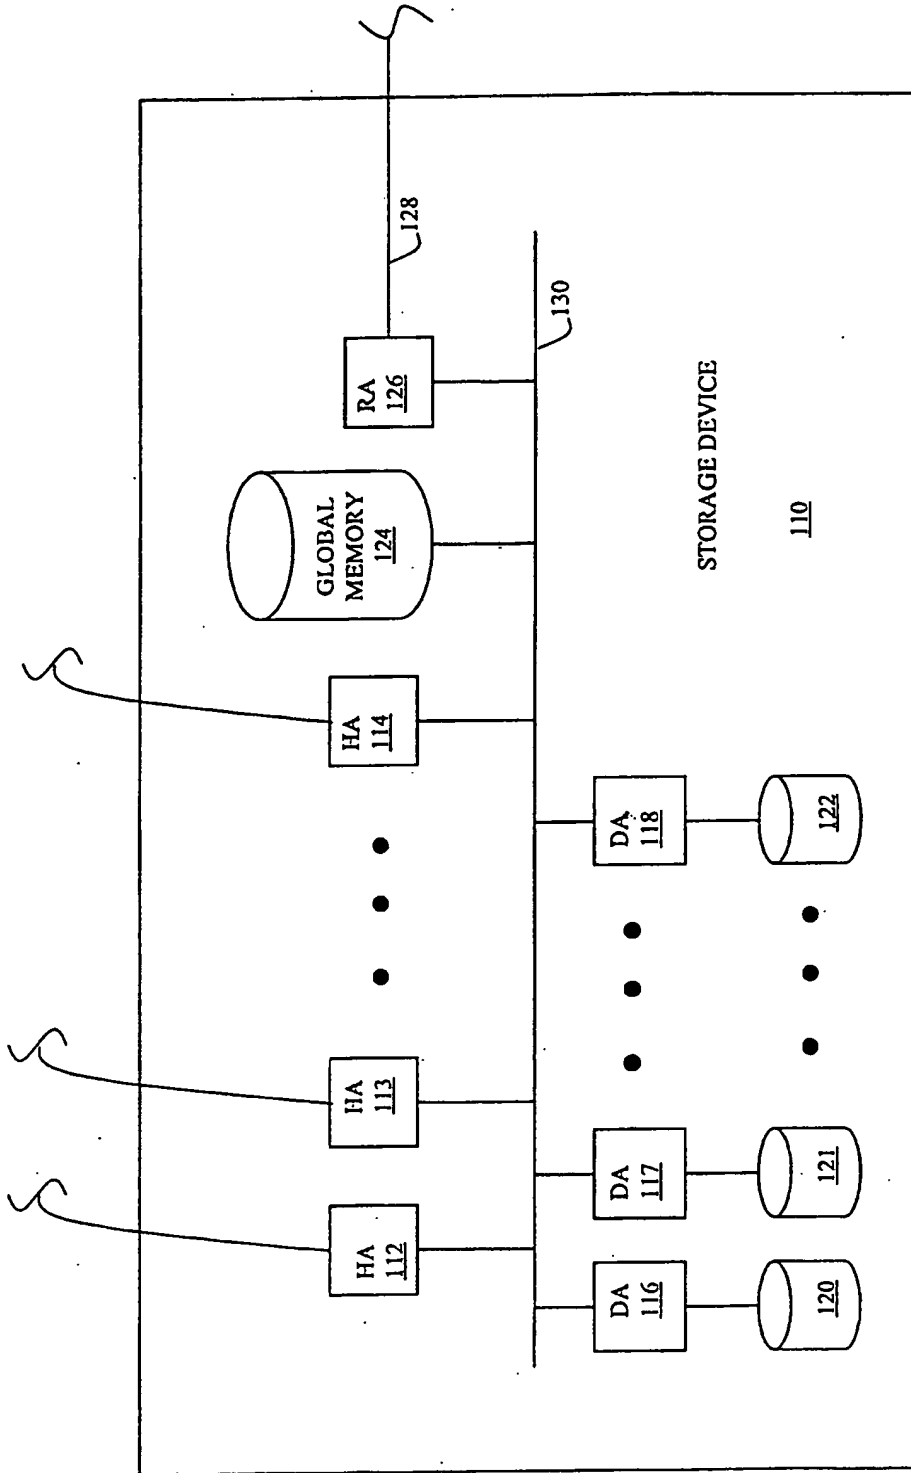


FIGURE 6

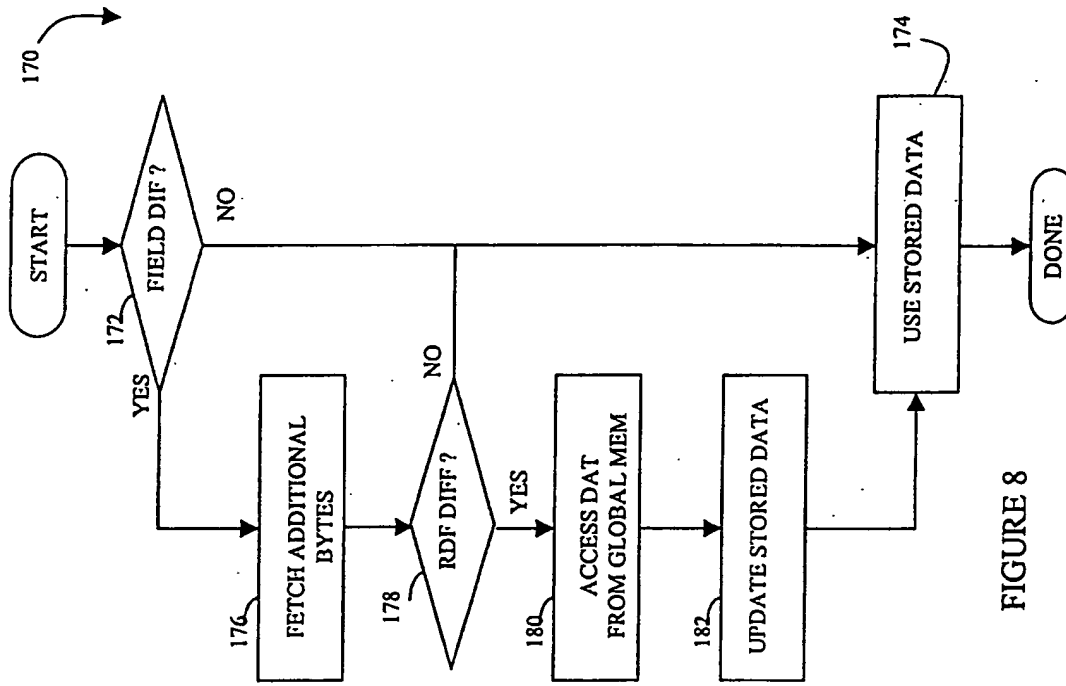


FIGURE 8

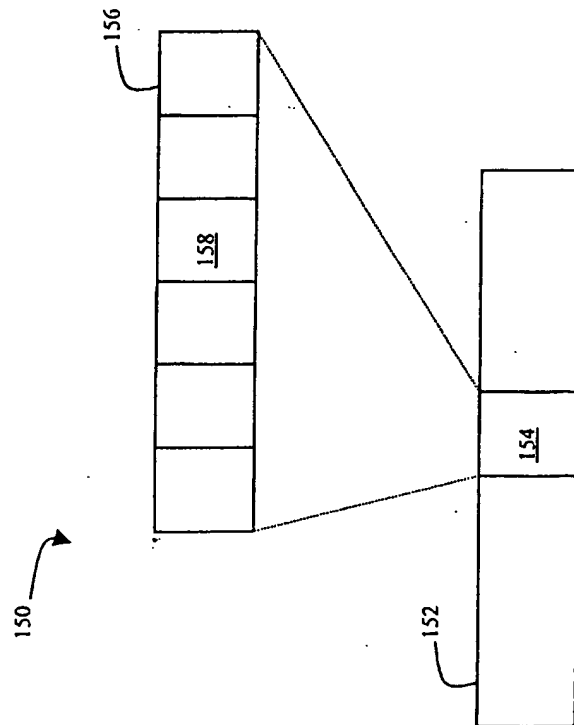


FIGURE 7

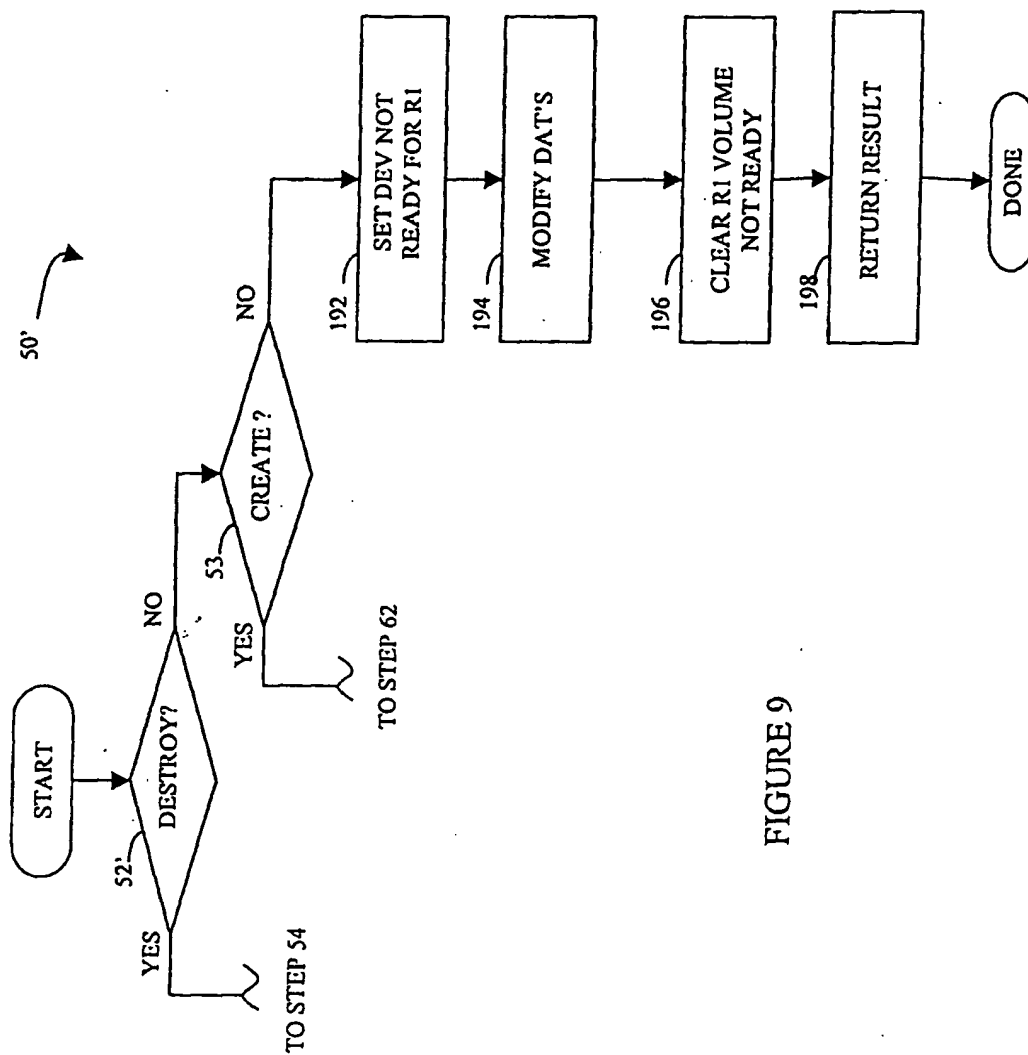


FIGURE 9



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 02 02 4809

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	US 5 895 493 A (GATICA JUAN C) 20 April 1999 (1999-04-20) * the whole document *	1-27	G06F11/20 G06F3/06
A	US 6 295 575 B1 (BLUMENAU STEVEN M ET AL) 25 September 2001 (2001-09-25) * column 28, line 27 - column 31, line 56 * * column 33, line 39 - column 34, line 52 *	1-27	
A	EP 1 093 051 A (EMC CORP) 18 April 2001 (2001-04-18) * the whole document *	1-27	
A	US 5 799 323 A (MOSHER JR MALCOLM ET AL) 25 August 1998 (1998-08-25) * the whole document *	1-27	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			G06F
The present search report has been drawn up for all claims			
Place of search MUNICH		Date of completion of the search 27 March 2003	Examiner Bozas, I
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.02 (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 02 02 4809

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

27-03-2003

Patent document cited in search report		Publication date		Patent family member(s)	Publication date
US 5895493	A	20-04-1999	AU	8272298 A	19-01-1999
			DE	69801824 D1	31-10-2001
			DE	69801824 T2	28-03-2002
			EP	0993632 A1	19-04-2000
			JP	2001512255 T	21-08-2001
			WO	9900727 A1	07-01-1999

US 6295575	B1	25-09-2001	US	2002007445 A1	17-01-2002

EP 1093051	A	18-04-2001	US	6341333 B1	22-01-2002
			EP	1093051 A2	18-04-2001
			JP	2002032197 A	31-01-2002

US 5799323	A	25-08-1998	US	5835915 A	10-11-1998
			US	5740433 A	14-04-1998
			US	5799322 A	25-08-1998
			US	5794252 A	11-08-1998
			CA	2167902 A1	25-07-1996
			DE	69614009 D1	30-08-2001
			DE	69614009 T2	15-11-2001
			EP	0724223 A1	31-07-1996

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82